

ASPIRE Supersonic Parachute Shape Reconstruction

Jason Rabinovitch*, Gregory S. Griffin†, William Seto‡, Clara O’Farrell§, Christopher L. Tanner¶, and Ian G. Clark||
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, 91109

This paper describes 3D stereo reconstruction efforts for the full-scale supersonic parachute deployment tests performed under the NASA Advanced Supersonic Parachute Inflation Research Experiment (ASPIRE) program. Specific time-resolved data products that have been generated for the ASPIRE tests and are presented in this work include: parachute projected area, parachute vent-hole location, and parachute surface maps (point clouds).

I. Introduction

Supersonic disc-gap band (DGB) parachutes have been used to successfully land assets on the surface of Mars since the 1970s [1–5]. The primary objective of the NASA Advanced Supersonic Parachute Inflation Research Experiment (ASPIRE) program is to demonstrate the capability of full-scale DGB parachutes in a Mars-relevant supersonic test environment. Three successful supersonic deployments of full-scale DGB parachutes have taken place as part of the ASPIRE test campaign. This paper will focus on the 3D canopy reconstruction efforts based on stereo image data acquired by the on-board science cameras that are focused on the parachute canopy during the flights.

II. ASPIRE Test Campaign Overview

The Advanced Supersonic Parachute Inflation Research Experiments (ASPIRE) project was started in 2016 to develop NASA’s capability to test supersonic parachutes at Mars-relevant conditions using a sounding rocket test platform. The initial series of ASPIRE flights were focused on testing two candidate parachute designs for the Mars 2020 project: a build-to-print version of the Mars Science Laboratory (MSL) disk-gap-band (DGB) parachute and a strengthened parachute with the same geometry but stronger materials [6]. The parachutes were delivered to targeted deployment conditions representative of flight at Mars by NASA’s Sounding Rocket Operations Contract (NSROC) sounding rockets. Three flights took place between October 2017 and September 2018. During the first test (SR01) [7], a build-to-print DGB was tested at conditions similar to those that MSL recorded during Mars entry, descent, and landing (EDL). During the second and third tests (SR02 and SR03, respectively), the strengthened parachute was tested at incrementally higher loads [8]. All three flights were successful, resulting in minimal damage to the parachute test articles.

The concept of operations for the ASPIRE sounding rocket tests is outlined in Fig. 1. The sounding rocket assembly, consisting of a Terrier first stage, a Black Brant second stage, and the roughly 1200 kg payload section containing the experiment, are launched out of NASA’s Wallops Flight Facility (WFF). The system is rail-launched and spin-stabilized at 4 Hz. During flight, the first and second stages burn out at altitudes of approximately 700 m and 16.7 km respectively, and the payload section reaches apogee between 50 km and 55 km. When the payload reaches the target dynamic pressure (q_∞) and Mach number conditions, the parachute is mortar-deployed. The deployment, inflation, and supersonic and subsonic aerodynamics of the parachute are analyzed by a suite of instruments including: a three-camera high-speed/high-resolution stereographic video system trained on the parachute, situational awareness video cameras, a set of load pins at the interface of the parachute triple-bridle and the payload, and a GPS and inertial measurement unit (IMU) onboard the payload. After decelerating to subsonic speed, the parachute and payload descend to the ocean for recovery and inspection.

Figure 2 shows a schematic of the ASPIRE configuration prior to launch from WFF. The payload section has a maximum diameter of 0.72 m, and is 7.54 m in length. It is comprised of a nose cone ballast section that is jettisoned before splashdown, a section housing electronics and foam for buoyancy, a telemetry section, a section housing

*Mechanical Engineer, Entry, Descent, & Landing and Formulation Group, 4800 Oak Grove Drive, M/S: T1708-112, Senior Member AIAA.

†Robotics Technologist, Aerial and Orbital Image Analysis Group, 4800 Oak Grove Drive, M/S: 198-219.

‡Robotics Technologist, Maritime and Multi-Agent Autonomy Group, 4800 Oak Grove Drive, M/S: 198-219.

§Guidance and Control Engineer, EDL Guidance & Control Systems Group, 4800 Oak Grove Drive, M/S: 321-220, Member AIAA.

¶Mechanical Engineer, Entry, Descent, & Landing and Formulation Group, 4800 Oak Grove Drive, M/S: 321-220, Member AIAA.

||Systems Engineer, EDL Systems & Advanced Technologies Group, 4800 Oak Grove Drive, M/S: 321-400, Member AIAA.




Fig. 1 ASPIRE concept of operations.

the attitude control system (ACS), and the experiment section containing the parachute mortar tube and ASPIRE instrumentation. The payload also includes an aft-transition section that remains attached to the Black Brant motor at separation. After separation from the Black Brant, the mass of the payload, including the stowed parachute and 600 kg nose cone, is approximately 1200 kg.

Figure 3 shows a schematic of the ASPIRE configuration after parachute deployment. The relevant dimensions of the parachute-payload system are labeled in the schematic, and their values are also listed in Table 1. The MSL build-to-print parachute tested on SR01 was an 80-gore DGB with a nominal diameter (D_0) of 21.35 m. The majority of the canopy is constructed using a Nylon fabric with a rated strength of 42 lbf/in. The parachute is built using a continuous line construction where the suspension lines continue in to the canopy as the radials cross the vent and continue along the opposite side of the canopy. The radials/suspension lines are constructed from 2100 lbf Technora line. The entire packed parachute assembly has a mass of 55.8 kg.

The strengthened parachute tested on SR02 and SR03 is an 80-gore DGB with a design nominal diameter (D_0) of 21.45 m. The geometry of the strengthened canopy is identical to that of the build-to-print DGB, but higher-strength materials are used in its construction. The broadcloth gores are constructed using a custom Nylon fabric with a rated strength of 110 lbf/in. The radials/suspension lines are constructed from 3200 lbf Technora line and the entire packed parachute assembly has a mass of 85 kg.

Table 2 shows the conditions at the beginning of parachute inflation for the three flights. Inflation is considered to start at suspension line stretch, i.e. when the entirety of the canopy exits the bag. Table 2 lists the wind-relative Mach number and dynamic pressure at line stretch. The time to full inflation is defined as the time from line stretch until the point when the parachute first achieves peak load. This first peak in parachute load may not be the highest peak, as supersonic parachutes often show a collapse and re-inflation. The last column in Table 2 shows the peak measured parachute load, regardless of when it occurred.

III. Experimental Data

The ASPIRE payload is equipped with scientific instrumentation to measure the trajectory, aerodynamics, and performance of the test vehicle and test article. A summary of the onboard instrumentation is provided in Table 3. The trajectory of the payload during each test is reconstructed using the measurements from the GLN-MAC IMU, the onboard GPS receiver, and three ground-based radars by means of an extended Kalman filter [9, 10]. In addition, meteorological balloons carrying radiosondes were released hourly during launch operations. These provide measurements of the

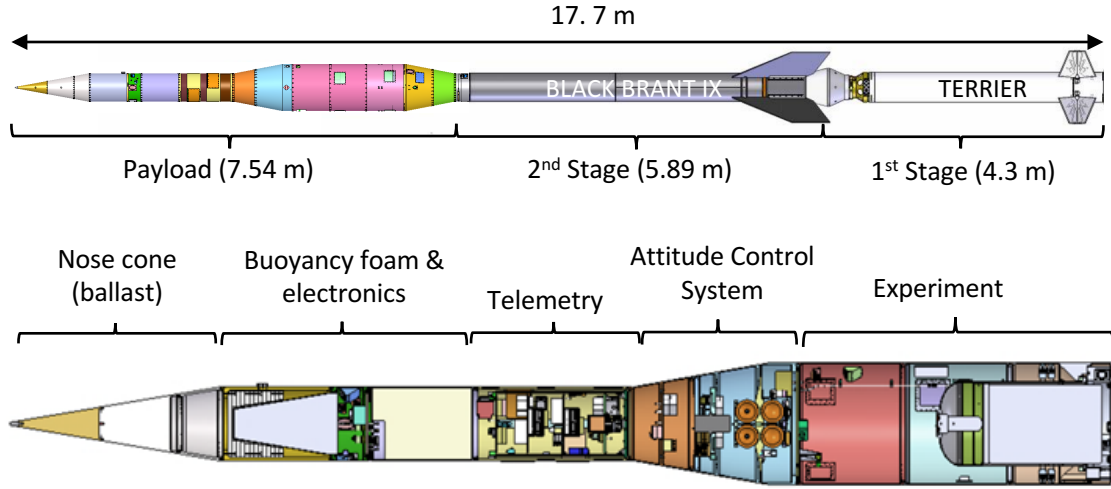


Fig. 2 Schematic of the ASPIRE launch configuration and payload.




Fig. 3 ASPIRE parachute system.

Table 1 As-designed dimensions of the ASPIRE parachutes.

Item	Symbol	SR01	SR02 & SR03
Parachute reference diameter (m)	D_0	21.35	21.45
Parachute reference area (m ²)	S_0	358.00	361.36
Vent diameter (m)	D_V	1.50	1.50
Disk diameter (m)	D_D	15.44	15.62
Gap height (m)	H_G	0.90	0.9 0
Band height (m)	H_B	2.60	2.60
Geometric porosity	λ_g	12.5%	12.8%
Suspension line length (m)	L_S	36.47	36.47
Riser Length (m)	L_R	7.06	7.78
Bridle Length (m)	L_B	1.35	1.30
Forebody diameter (m)	d	0.72	0.72

Table 2 Parachute deployment conditions.

Test	Mach at line stretch	q_∞ at line stretch	Inflation time	Peak load	Density at line stretch	Static pressure at line stretch
SR01	1.79	492 Pa	0.506 sec	32.4 klbf	3.05×10^{-3} kg/m ³	218 Pa
SR02	2.00	745 Pa	0.456 sec	55.8 klbf	3.68×10^{-3} kg/m ³	266 Pa
SR03	1.88	1028 Pa	0.410 sec	67.4 klbf	6.02×10^{-3} kg/m ³	415 Pa

temperature, density, and winds up to 40 km.


Table 3 Key instrumentation summary

Device	Sample Rate / Resolution	Notes
GLN-MAC IMU	400 Hz	Provides both 100 Hz and 400 Hz raw data
Javad TR-G2 GPS	20 Hz	
Radars (x3)	50 Hz	
Load Pins [†]	1 kHz	90 klbf rated, calibrated to 30 or 40 klbf*
HS Cameras (x3) [†]	1000 fps, 3840x2400	Global shutter, 12-bit grayscale images
Situational Video (x2) [†]	120 fps, 1920x1080	Rolling shutter
Situational Video (x1) [†]	30 fps, 3840x2160	SR02 and SR03 only. Rolling shutter
Radiosondes (x4)	1 Hz	Balloon-launched every hour

[†] See Fig. 4.

* Calibrated to 30 klbf on SR01 and SR02, and to 40 klbf on SR03.

Figure 4 shows the aft-most portion of the payload, which houses the parachute mortar. A specially-designed instrumentation ring, also shown in Fig. 4, fits around and above the parachute mortar tube, and houses many of the key instruments. Three load pins are installed on the ring, at the locations where the parachute triple bridge legs attach to the payload. This allows measurement of the aerodynamic forces on the parachute. The load pins are rated to 90 klbf, but calibrated to peak loads of 30 klbf (SR01 and SR02), or 40 klbf (SR03) to reduce the uncertainty in the load pin measurements.

**Fig. 4 Instrumentation ring and sensor locations**

Two sets of cameras are used to record the parachute deployment and inflation events: a set of high-resolution, high-speed (HS) cameras, and a set of situational awareness cameras. The HS camera system consists of three IDT OS-10 cameras paired with a ruggedized Schneider APO-Xenoplan 2.0/20mm compact C-mount lens. The cameras are




Fig. 5 Challenges for 3D reconstruction of the parachute surface include: 1) occlusion from suspension lines and other support hardware, 2) complex topology during deployment with surfaces that remain nearly perpendicular to the camera line of sight (especially near the canopy skirt), 3) uncontrolled lighting conditions leading to shadows and non-Lambertian reflections, and 4a) variations in pixel gain over the sensor, shown here before and 4b) after flat-field gain correction.

mounted around the circumference of the instrumentation ring (which has a radius of 29.6 cm at the camera mounting locations), at 120° intervals. They record 12-bit grayscale images at 1000 frames per second at a 4K resolution. Recording is performed with an aperture of 4.0 and a $40 \mu\text{s}$ exposure time.

The situational awareness cameras are a set of standard GoPro Hero 4 cameras. The GoPros provide a wider angle contextual view and an NTSC feed that is telemetered to ground during the flights. Two of these cameras are set to record at 1080p and 120 fps, and to transmit video over NTSC during flight. On SR02 and SR03, a third GoPro camera was set to record 4K video at 30 fps.

IV. Image Processing

This section provides an overview of the image processing techniques used to produce 3D point-clouds representing the surface of the parachute canopy. These point-clouds are generated from the high-speed images taken during the ASPIRE flights.

A. Overview

Figure 5 illustrates some of the difficulties encountered when using the three HS cameras for 3D reconstruction of the parachute surface during flight. The parachute surface, which lacks fiducials or rigid elements for distance calibration, twists and folds at angles that can be nearly perpendicular to the camera line-of-sight. Uncontrolled lighting conditions during flight lead to dark regions and specular reflections which further complicate the analysis [11]. Small gain variations in individual pixels must be corrected to accurately establish correspondences between images.

Out of the various factors that limit surface coverage and resolution, three were found to be particularly significant:

- 1) Occlusions: fast-moving suspension lines and other support structures obstruct the view of the parachute surface during the flights. These foreground structures can be particularly confusing for many popular dense stereo reconstruction methods that rely on template-based disparity searches over epipolar lines in rectified image pairs.
- 2) Camera movement: despite efforts to stabilize the camera mounts, it is believed that high g-forces lead to small but significant changes in camera pose that must be corrected during image post-processing.
- 3) Camera baselines: three HS cameras equally spaced around the instrumentation ring form three baselines with a length of 0.514 m. The surface reconstruction error varies inversely with this baseline distance [12], which is constrained by the radius of the sounding rocket.

In an indoor wind-tunnel environment under controlled lighting conditions, camera spacing can be increased as needed along stable mounts located *behind* the parachute, where suspension lines and other obstructions are not

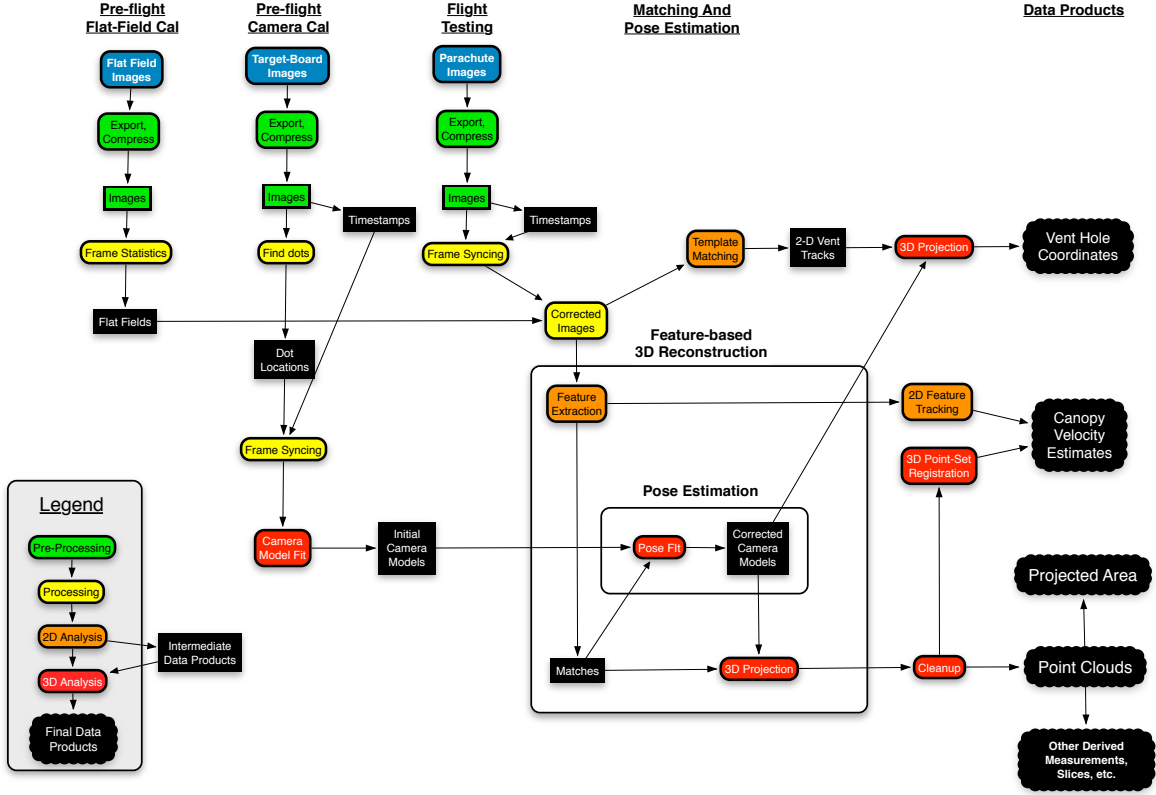


Fig. 6 Image analysis workflow.

Table 4 Results of pre-flight camera calibration for the three HS cameras used in SR01, SR02 and SR03. Shown here are 2 intrinsic and 2 extrinsic parameters: horizontal and vertical field of view (hFOV and vFOV, respectively), and relative distance and orientation relative to camera 1.

	HS Cam1		HS Cam2				HS Cam3			
	hFOV	vFOV	hFOV	vFOV	d_{12}	θ_{12}	hFOV	vFOV	d_{13}	θ_{13}
SR01	47.7°	31.4°	47.7°	31.5°	.514 m	60.1°	47.7°	31.4°	.514 m	-60.0°
SR02	47.9°	31.6°	47.8°	31.6°	.513 m	60.0°	47.9°	31.6°	.513 m	-60.2°
SR03	47.8°	31.4°	47.8°	31.5°	.514 m	60.2°	47.9°	31.5°	.514 m	-60.1°

visible [13]. For the purpose of the current in-flight surface reconstruction efforts, we instead adopt an analysis workflow that minimizes the sources of error that cannot be controlled.

As shown in Fig. 6, the workflow begins with a pre-flight flat-field calibration. By placing grey and white cards of uniform brightness in front of each camera lens, the resulting out-of-focus images are used to establish a map of gain fluctuations over individual sensor pixels. This correction significantly increases the number and quality of correspondences available for pose correct and point-cloud reconstruction, as described in Sec. IV.C. For both the pre-flight and in-flight imagery, GPS timestamps are used to re-sequence the frames to correct for a small (< 1 per 3000) fraction of dropped images. The resulting image sets are synchronized to within $13\mu\text{s}$, $25\mu\text{s}$ and $16\mu\text{s}$ for SR01, SR02, and SR03, respectively.

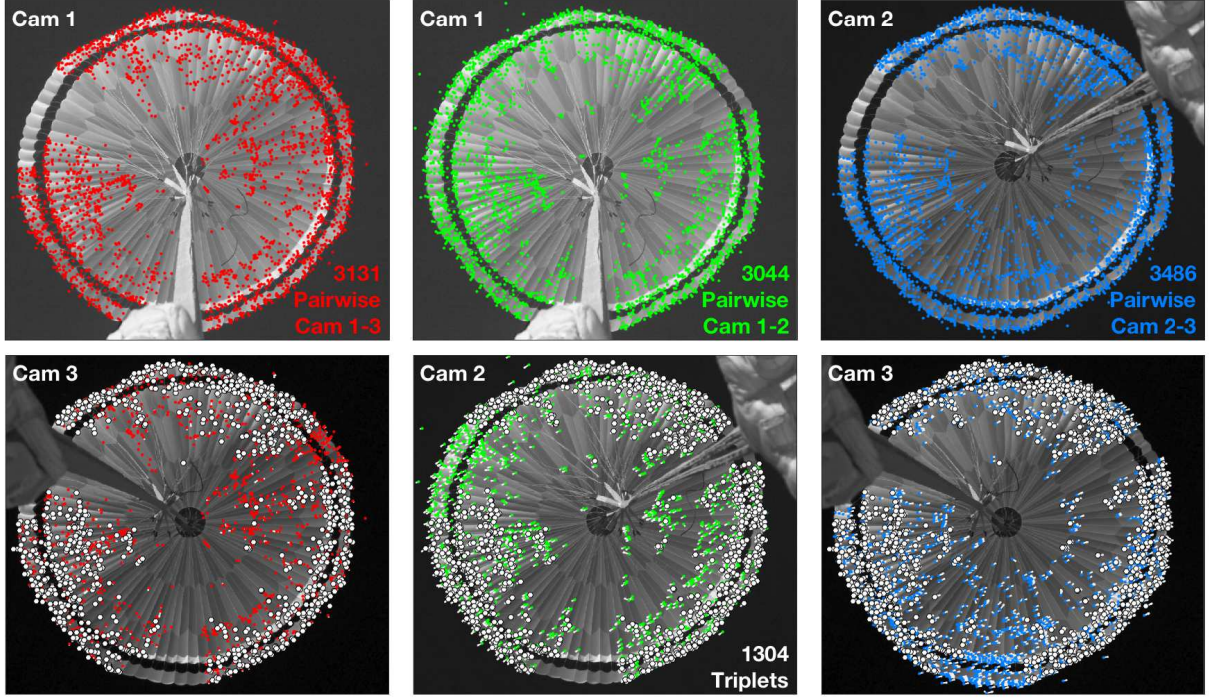


Fig. 7 Example showing SIFT feature correspondences matched across a single set of synchronized images taken 1.500s after mortar fire. Typically, about 2500-3500 pairwise matches were found for each image pair, resulting in 1000-1500 triplets over each set of 3 images. These matches are used for both camera pose-estimation and for generating a high-speed 3D reconstruction of the parachute canopy as it is deployed.

B. Camera Calibration

Pre-flight camera calibration tests are performed near the launch facility. With the instrumentation ring pointed towards the horizon, a flat target board with regularly-spaced dots is carried across the overlapping field of view of the HS cameras. Care is taken to image these boards at a variety of locations and distances, covering as much of the focal plane of each sensor as possible. A common analysis procedure is used whereby dot positions are extracted in each image and triangulated across N sets of image triplets to produce 3-D coordinates in space. These coordinates, in turn, are reprojected back onto the individual image planes for comparison with the original dot positions [14–17]. A least-squares fit procedure optimizes the camera intrinsic and extrinsic parameters to minimize the overall reprojection error in the image plane. Because each 10×10 array of dots is constrained to lie along a regular planar grid, the degrees of freedom in the fit are drastically reduced, from $300N$ dot coordinates to $6N$ coordinates representing the individual board positions and orientations.

The specific implementation in this work is performed in MATLAB and uses CAHVOR models [18, 19] with 8 intrinsic parameters per camera, and 6 additional extrinsic parameters specifying the positions and orientations of cameras 2 and 3 with respect to camera 1. Thus a total of 36 model parameters are minimized using dots extracted from 80 – 120 images of target boards placed at distances between 3 and 60 meters. The final mean re-projection error over all dots in all images varied from 0.29 to 0.34 pixels in the different calibrations performed. Table 4 gives the resulting field of view and the relative distance and orientation of the three cameras.

C. Pose Estimation and Point Cloud Reconstruction

This work implements a method that uses matched feature keypoints to correct camera extrinsics, while simultaneously reconstructing the parachute canopy. Figure 7 illustrates the process, which starts by extracting and matching SIFT keypoint and features [20] between image pairs using the MATLAB toolbox in VLFEAT [21]. The problem of using image correspondences to estimate pose while constructing detailed point clouds has been well studied [22]. The current software is adapted from routines developed by Ansar et al. [23] for use in robotics applications at the Jet Propulsion

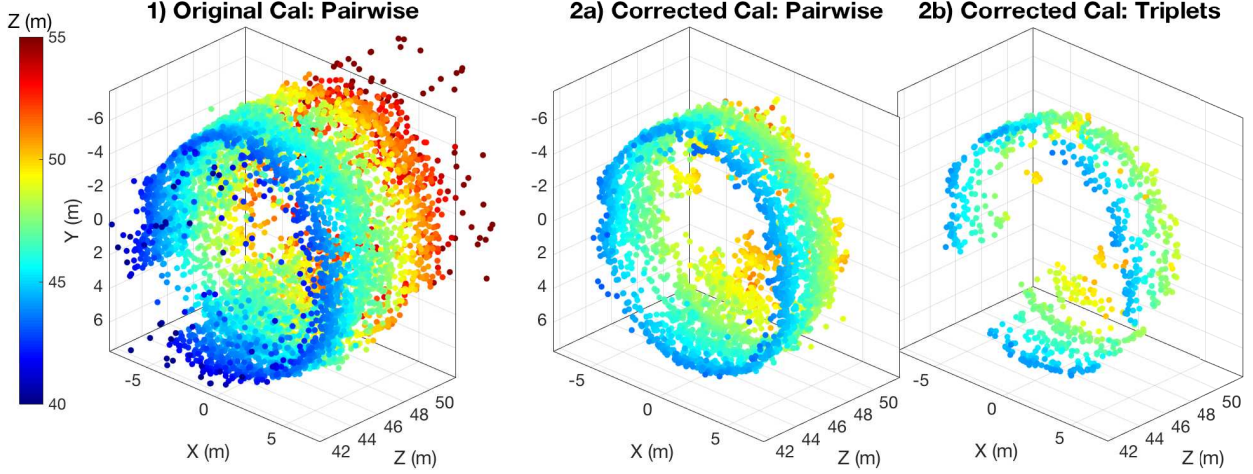


Fig. 8 Point clouds generated by triangulating the feature correspondences shown in Fig. 7. Each point cloud is color-coded by depth. In the left-most plot, pairwise matches - aggregated over all 3 baselines - are triangulated using camera models calibrated prior to launch. The distorted appearance seen here is a consequence of small changes in camera pose that occurred after the original camera calibration. By minimizing the reprojection error of feature triplets, the pose is corrected resulting in 2a) a pairwise surface reconstruction along with 2b) a more accurate but less complete reconstruction based on feature triplets.

Laboratory, California Institute of Technology [24, 25].

As in Sec. IV.B, camera pose is corrected using a least-squares solver that minimizes reprojection error. Instead of using uniform grids of dots on planar surfaces, the pose correction step uses feature triplets found anywhere in the overlapping field of view of the 3 cameras. During flight, most of these triplets will be correspondences discovered on the canopy surface. The objective function triangulates these triplets and reprojects their intersection points back onto the image planes. The pose is optimal when it minimizes the residual error between the original and reprojected triplets.

Unlike target-board dots, feature matches can occur anywhere in the scene at unspecified locations. Despite this lack of constraints, the minimization procedure still works in practice because of 1) the large number of points available for fitting, 2) the high quality of matches, and 3) the small number of fit parameters. Specifically, we vary only 6 parameters that are believed to be the most likely to change during flight, namely: the pitch, yaw and roll of cameras 2 and 3 with respect to camera 1. High quality matches are obtained by enforcing a consistency check over the 3 images: three pairwise matches form a “triplet” only if successive matches from camera 1 to camera 2, camera 2 to camera 3, and camera 3 to camera 1 lead back to the exact same coordinate. A missing feature or an erroneous association between one or more of these image pairs is likely to cause this check to fail. In practice, less than half of pairwise matches are part of a triplet.

As the parachute canopy is visible for at least 4.8s in each 1000 Hz image collection, roughly 5 million triplets are extracted during flight. Of these, 100,000 triplets are randomly chosen to generate a single set of corrected camera models for the flight. The magnitude of the final corrections in pitch, yaw and roll are on the order of a few tenths of a degree, with the smallest changes occurring in roll. Figure 8 illustrates the importance of correcting the original camera models for small changes in pose. Before correction, poor alignment between separate camera pairs results in the large depth disparity observed in the left-most plot. By comparison, the corrected point clouds on the right show noticeably less surface variation.

To track the evolution of the parachute canopy over time, two separate point clouds are generated at 1 ms time intervals throughout the flight. These are the *pairwise* and *triplet*-based point clouds plotted on the right side of Fig. 8.

Triplet-based Clouds Because depth uncertainty typically decreases when the same point is viewed from multiple angles [26], the point cloud based on triplets is intended to provide the most accurate representation of the parachute canopy. This accuracy comes at the expense of coverage, however, since a features that is occluded in any 1 of the 3 camera views will not be part of a triplet.

Pairwise-based Clouds For applications where coverage is more important than accuracy, such as the projected area calculations in Sec. V.B, a second point cloud based on pairwise matches provides more complete coverage.




Fig. 9 The trajectory of the central vent hole as calculated for SR02. The location is tracked using a disk-shaped template that matches the vent hole. As the size and appearance of the vent hole does not change significantly, it can be followed without specialized fiducial markers despite intermittent occlusion from suspension lines and other hardware. Here the trajectory is plotted in the Camera Coordinate System, with its origin at the midpoint between 3 HS cameras oriented along the z-axis.

This is especially true near the middle of the canopy, where support lines converge and cause the greatest occlusion and confusion.

Based on localized gradient histograms, the SIFT features used in our analysis have a number of properties that help to address challenges discussed in Sec. IV.A. They provide limited invariance to lighting variations, affine transformations, and small changes in the appearance of surface features as seen from different camera angles. They are, however, not immune to occlusion from suspension lines and other parachute support structures, which can lead to spurious false matches. For this reason, the final point clouds generated in Fig. 6 are “cleaned” by back-projecting each point and computing its error with respect to the original matches. Typically 2-5% of points with errors greater than 0.5 pixels are removed from each cloud.

V. Data Products

This section will outline the different data products that are calculated based on the 3D points clouds generated during the ASPIRE tests. Detailed data processing is performed for the SR02 and SR03 flights only, as only two of the three science cameras recorded data during the SR01 flight, whereas all three cameras recorded data for SR02 and SR03.

A. Vent Hole Tracking

The vent hole located in the middle of the parachute is tracked using a disk-shaped template that matches the size and shape of the vent hole. Coordinates are estimated by 1) convolving each image with the template, and 2) interpolating the location of the maximum in the resulting convolved images [27]. These coordinates are then compared across synchronized image sets to triangulate the 3D trajectory of the vent hole during flight. The results of this process are shown in Fig. 9. Approximately 8-12% of coordinates are removed based on low match scores or high reprojection

error, caused mainly by intermittent obstruction of the hole by suspension lines in one or more camera views.

The largest error in the trajectory is observed in the depth direction, where small errors in the estimated image coordinates lead to depth errors that scale quadratically with distance [12]. For this reason, two vent-hole data products are generated: 1) the original triangulated coordinates, and 2) a corresponding set of distance-smoothed coordinates. The smoothing process involves converting the trajectory to spherical coordinates and replacing $r(t)$ with a best-fit 3rd-order polynomial. This step removes transient errors in $r(t)$ that are likely to be artifacts of intermittent occlusion.

B. Projected Area Calculations

To calculate the projected area of the parachute, a computed point cloud based on pairwise matches with coordinates in the camera frame is used. To project the points into a 2D cross-section, the Z-coordinate is dropped from the 3D points, such that the area will be determined in a plane perpendicular to the line of sight of the cameras. The boundary of the 2D points is computed, and then the area is calculated from these ordered boundary points. To determine the boundary of the parachute, a convex hull of the points is computed. It is observed that the largest cross-section of the parachute does not exhibit much concavity throughout flight, and the convex hull allows allows a relevant area to be calculated even if the entire outer edge of the parachute is not resolved during the 3D reconstruction. Figure 10 illustrates the convex hull for a set of points, as well as a tighter boundary generated from a naive boundary fitting algorithm. If there are occluded regions in the point cloud, a more aggressive choice of the “shrink factor” in the algorithm can cause the algorithm to underestimate the projected area of the parachute, as the algorithm has no prior information or constraints on the shape of the parachute boundary. Although future work will involve estimating a tighter boundary of the of the parachute with respect to the border generated by the convex hull, in practice, the area generated using the convex hull boundary differs from the area calculated with a tighter boundary only by approximately 1-2%.

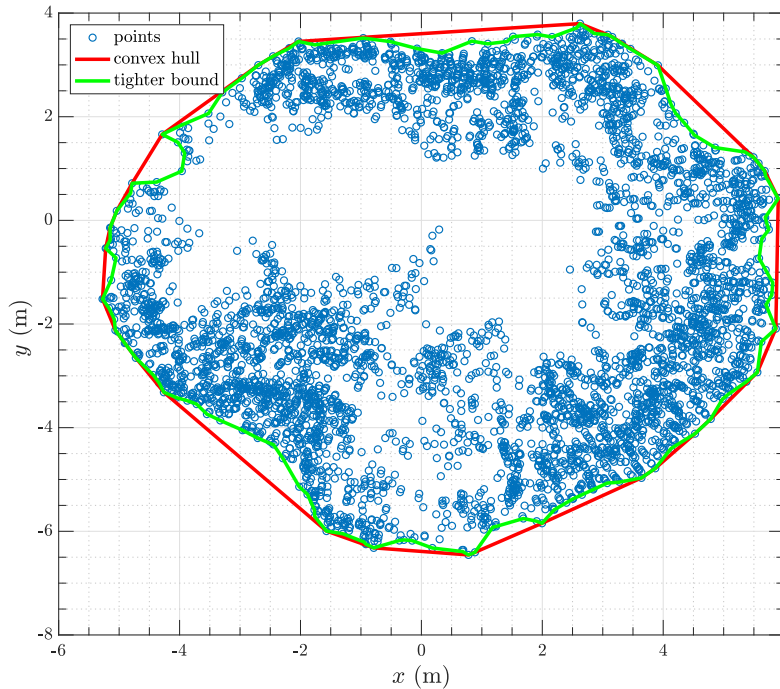


Fig. 10 Sample illustration of projected area calculation where two different algorithms are used to calculate the parachute boundary.




Fig. 11 Tilt Angle (from vent hole) compared with area for SR02.

1. Error sources

Since the method of computing area based on a polygon composed of 2D points is fairly straightforward, the error of this method can be estimated based on the uncertainty of the chosen boundary points. To perform this error estimation, the initial uncertainty for the XY coordinate of the point, currently $\sigma_{point} = 0.5$ m, is propagated through the Shoelace formula, which is the equation used to compute the area of the polygon formed by the points. This initial point uncertainty is estimated empirically in a separate procedure by selecting some points on the boundary of the canopy, sampling a large number of noisy points from a normal distribution with the image coordinate as mean and standard deviation representing calibration and matching error, and then triangulating these noisy points and observing the standard deviation. The final uncertainty for the area calculation is $\sigma_{area} \approx 1.5$ m², which can be as high as $\sim 15\%$ relative error for times early in the inflation process, but is $\sim 1\%$ after full inflation.

While the baseline uncertainty from the error in triangulation of 3D points is small, it is believed that the main contribution to error when estimating the projected area is whether the boundary points that are selected accurately represent the largest cross-section of the parachute. The primary source of error that affects whether these points can be captured are occlusions of the parachute skirt. While one major source of occlusion is the confluence fitting and riser, this is mostly mitigated by having a 3rd camera, which fills in the occlusions of the two other cameras. Another potential source of occlusion of the parachute is self-occlusion caused by large solid-body rotation of the parachute. Overlaying the projected area with the tilt angle computed from the vent hole location for SR02 (Fig. 11), there appears to be a correlation between a larger rotation angle of the parachute and a drop in projected area. This drop in area is believed to be due to occlusion of aspects of the canopy, as the analytical decrease in the area due to the rotation does not account for the large drops observed, which are observed to be up to 8%. A possible reason as to why this effect is observed is that rotation of the parachute can cause the widest part of the skirt to be occluded, which prevents the reconstruction algorithm from identifying this part of the parachute. This phenomena is observed for SR02, where there is a relatively large parachute motion, but not for SR03 which has minimal motion, which is believed to contribute to the relatively smooth (fewer oscillations) calculated projected area throughout the flight. While it is possible to project the reconstructed points onto a plane that is perpendicular to a vector running from the cameras to the vent hole, this is not performed as it does not remove the error associated with camera occlusions, and would only account for a possible

solid-body rotation of the parachute.

C. Parachute Canopy Velocity

Two methods are used to compute and compare canopy velocity measurements: an image-based method using feature points, and a method using the 3D coordinates of point clouds directly. For both methods, the points used are pre-thresholded by distance to attempt to capture points on the parachute skirt only. Figure 12 illustrates the results of computing velocities for points of interest from the two different methods for the same image frame.

In the image-based approach, feature points that were used to generate the point cloud are selected to be tracked across successive frames. Each frame is tracked independently, as each frame has an independent set of feature points. The points are tracked using the Kanade-Lucas-Tomasi (KLT) feature tracking algorithm [28]. After tracking, a pixel motion is obtained for each point, which is then rescaled into a dimensional velocity (m/s) using the following:

$$v_{dimensional} = \frac{v_{pixel} * Z}{f} * frame_rate, \quad (1)$$

where v_{pixel} is measured in pixels per frame, Z is the depth of the feature point in meters, f is the focal length of the camera in pixels, and $frame_rate$ is the frame rate, which is equal to 1000 frames per second for this work. The main benefit of this method is that estimating the motion in image space is accurate. However, this method is challenging to apply when large motions are observed from frame to frame, where large parachute motion and deformations will inhibit feature tracking.

In the direct approach which utilizes the 3D reconstructed points, the goal is to compute a motion for each point that transforms the current point cloud into the next point cloud. The point-wise motion will correspond to the velocity of the canopy points. The method of Coherent Point Drift [29] is employed for this approach, which performs non-rigid alignment of unorganized point clouds. This method represents the point sets using a Gaussian mixture model. The benefit of representing the data with a Gaussian mixture model is that the correspondence problem between point clouds does not have to be solved, as noisy data is handled naturally, and the computed motion is smooth. Points that are near each other should have similar velocities. In contrast, with the image-based method, the measurements generated from this point set alignment may be less accurate because the inputs are noisy triangulated points and contain outliers. Image feature tracking generally only works well with small motions and the assumption of constant brightness of the feature points, which can limit this method to small frame differences for the current application. However, the point set alignment method can potentially compute motion across 10 s of frames since the inputs are simply point clouds, as long as the shapes are not dramatically different. Future work could involve smoothing the results over a larger range to overcome the single frame inaccuracies. Points are aligned in 2D to decrease processing time and focus on velocity components in the X-Y plane.


As an optional post-processing step, results are also generated with the vent hole velocity subtracted. This is useful for the post-inflation phase, where large solid-body parachute motion can dominate the observed motion. Figure 13 illustrates an instance in the trajectory where the dominant motion is a rigid body mode by the entire canopy, as opposed to local relative motion. After subtracting the vent velocity, the actual velocity of the feature points is nearly zero.

VI. Discussion


While the majority of the current effort has focused on generating the previously described data products, some preliminary work has been completed on analyzing the data products to gain better insight into the supersonic parachute dynamics observed during the ASPIRE flight tests.

A. Load and Projected Area Data

The calculated projected areas of the parachutes for SR02 and SR03 are shown in Fig. 14, and the comparison of the normalized projected area and measured load is shown in Fig. 15. A general trend seen in the projected area data is that the projected area tends to lag behind the peaks in measured load. Furthermore, as the measured load decreases post-inflation for both SR02 and SR03, a small increase in projected area is observed.




(a)




(b)

Fig. 12 Two different methods of computing canopy velocity, a) image feature tracking then rescaled to calculate velocity, and b) point set alignment.



(a)



(b)

Fig. 13 Canopy velocity without (a) and with (b) subtracting the vent hole smoothed velocity.




Fig. 14 Projected area vs time for SR02 and SR03 flights.






Fig. 15 Normalized load and projected area for SR02 and SR03 flights. The load and projected area are normalized by each of their respective maximum values.

B. Inflation Dynamics

The same numerical method for point alignment which was used to compute the parachute canopy velocity can also be leveraged to fit a contour to desired points in a 3D point cloud. Fitting contours to the points which comprise the skirt of the parachute can provide some insight into the dynamics that occur during canopy inflation. Figure 16 shows a calculated contour that has been fitted points found to be in the parachute skirt, and then the points are re-projected back onto the original image in order to visualize how accurately the skirt is captured. Figure 17 shows fitted canopy skirt contours color-coded by time for SR02 and SR03. The times shown span the inflation process up to the point where peak load is measured. From Sec. V.C, it can be seen that parts of the canopy skirt are moving at ~ 100 m/s during inflation.




(a)




(b)

Fig. 16 Estimated skirt contour overlaid on the original parachute image frame for SR02 (a) and SR03 (b).



(a)



(b)

Fig. 17 Fitted contours to the skirt points color-coded by time from line stretch for (a) SR02, and (b) SR03.

C. Inflated Canopy Shape

Figure 18 shows two sample point clouds from the SR02 flight test. As previously discussed, the discrepancy between the point cloud generated with pairwise matches vs triplets is clearly seen, with the triplet matches resolving less of the canopy surface. However, it is believed that the error associated with the XY coordinates of the parachute surface generated with the triplet matches are lower than the error associated with the points with pairwise matches.




Fig. 18 Sample point clouds from the SR02 flight. (a) Pairwise matches, and (b) triplet matches. The point clouds show the location of features on the canopy 1.89 sec after mortar fire. Features have been colored by their Z coordinate.

Figure 19 shows variations of the SR03 canopy profile post-inflation. Colors correspond to the probability that the reconstructed canopy shape corresponded to a specific profile. In addition, the experimental data is compared to an FEA derived symmetric DGB profile. Contours in Fig. 19 represent a normalized PDF of Z-coordinates of the canopy location for a given r , where r is defined to be $r = ((X - x_c)^2 + (Y - y_c)^2)^{0.5}$, where x_c and y_c are the X and Y locations of the center of the vent hole (for a given frame). This allows a side profile view of the canopy to be visualized, similar to profiles presented in [13]. These distributions of Z-coordinates are created by using all triplet matches in images starting ~ 1 s after parachute deployment. Distributions are normalized on a per r -value basis in order to account for the fact that the triplet point clouds generally find more points near the skirt compared to the center of the parachute, as previously discussed. Without normalization, the structure of the canopy closer to the vent hole is more difficult to visualize.

Figure 19 also includes an FEA derived DGB profile as a point of comparison. In one case the mean vent hole Z coordinate is used to align this profile with the experimental results, while the other curve uses the same profile but arbitrarily shifts the Z-coordinate for a better qualitative alignment with the experimental data. While the experimental data is quite noisy, there do seem to be some curvature differences between the experimental data and the FEA derived DGB profile.




Fig. 19 Reconstructed canopy shape PDF for SR03 (contours) compared to a FEA derived DGB profile (markers).

VII. Conclusion

This work has provided an overview of the ASPIRE test campaign and image processing techniques used to create time-resolved point clouds representing the parachute surface as a function of time. These point clouds were then used to further look into different physical phenomena associated with supersonic parachute dynamics. Due to a variety of error sources presented in this work the 3D reconstructed point clouds can be noisy, and preliminary work has been completed to compare the experimentally observed canopy shape with FEA derived axi-symmetric models.

Acknowledgments

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The authors would like to acknowledge the contributions of the entire ASPIRE team, and would also like to thank Yang Cheng and Adnan Ansar for providing software and advice.

References

- [1] Adams, D. S., Witkowski, A., and Kandis, M., “Phoenix Mars Scout parachute flight behavior and observations,” *Aerospace Conference, 2011 IEEEAC paper 1534*, IEEE, 2011. [1](#)
- [2] Cooley, C., and Lewis, J., “Viking 75 Project: Viking Lander System Primary Mission Performance Report,” Tech. rep., Martin Marietta Corporation, NASA/CR-145148, April 1977.
- [3] Spencer, D. A., Blanchard, R. C., Braun, R. D., Kallemeyn, P. H., and Thurman, S. W., “Mars Pathfinder Entry, Descent, and Landing Reconstruction,” *Journal of Spacecraft and Rockets*, Vol. 36, No. 3, 1999. doi:10.2514/2.3478, URL <http://dx.doi.org/10.2514/2.3478>.
- [4] Steltzner, A., Kipp, D., Chen, A., Burkhart, D., Guernsey, C., Mendeck, G., Mitcheltree, R., Powell, R., Rivellini, T., San Martin, M., et al., “Mars Science Laboratory entry, descent, and landing system,” *Aerospace Conference, 2006 IEEEAC paper 1497*, IEEE, 2006.
- [5] Witkowski, A., and Bruno, R., “Mars exploration rover parachute decelerator system program overview,” *17th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, 2003. [1](#)
- [6] Tanner, C. L., Clark, I. G., and Chen, A., “Overview of the Mars 2020 Parachute Risk Reduction Plan,” *IEEEAC Paper*, 2018. [1](#)
- [7] O’Farrell, C., Karlgaard, C., Tynis, J. A., and Clark, I. G., “Overview and Reconstruction of the ASPIRE Project’s SR01 Supersonic Parachute Test,” *IEEEAC Paper*, 2018. [1](#)
- [8] O’Farrell, C., Sonneveldt, B. S., Karlgaard, C., Tynis, J. A., and Clark, I. G., “Overview of the ASPIRE Project’s Supersonic Flight Tests of a Strengthened DGB Parachute,” *IEEEAC Paper*, 2019. [1](#)
- [9] Karlgaard, C., Tynis, J. A., and O’Farrell, C., “Reconstruction of the Advanced Supersonic Parachute Inflation Research Experiment Sounding Rocket Flight Test,” *AIAA Paper*, , No. 2018-3624, 2018. [2](#)
- [10] Karlgaard, C., Tynis, J. A., O’Farrell, C., and Sonneveldt, B., “Reconstruction of the Advanced Supersonic Parachute Inflation Research Experiment Sounding Rocket Flight Tests SR02 and SR03,” *AIAA Paper*, 2019. [2](#)
- [11] Wöhler, C., and d’Angelo, P., “Stereo Image Analysis of Non-Lambertian Surfaces,” *International Journal of Computer Vision*, Vol. 81, No. 2, 2008, pp. 172–190. [5](#)
- [12] Gallup, D., Frahm, J., Mordohai, P., and Pollefeys, M., “Variable baseline/resolution stereo,” *2008 IEEE Conference on Computer Vision and Pattern Recognition*, 2008, pp. 1–8. doi:10.1109/CVPR.2008.4587671. [5](#), [10](#)
- [13] Schairer, E. T., Kushner, L. K., Heineck, J. T., and Solis, E., “Measurements of Parachute Dynamics in the World’s Largest Wind Tunnel by Stereo Photogrammetry,” *2018 Aerodynamic Measurement Technology and Ground Testing Conference*, 2018, p. 3802. [6](#), [18](#)
- [14] Weng, J., Cohen, P., and Herniou, M., “Camera calibration with distortion models and accuracy evaluation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, Vol. 14, No. 10, 1992, pp. 965–980. [7](#)
- [15] Zhang, Z., “A flexible new technique for camera calibration,” *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, Vol. 22, No. 11, 2000, pp. 1330–1334.

- [16] Roth, G., "Computing camera positions from a multi-camera head," *Third International Conference on 3-D Digital Imaging and Modeling*, IEEE Comput. Soc, 2001, pp. 135–142.
- [17] Heikkila, J., and Silven, O., "A four-step camera calibration procedure with implicit image correction," *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 1997, pp. 1106–1112. [7](#)
- [18] Di, K., and Li, R., "CAHVOR camera model and its photogrammetric conversion for planetary applications - Di - 2004 - Journal of Geophysical Research: Planets (1991–2012) - Wiley Online Library," *Journal of Geophysical Research*, Vol. 109, No. E4, 2004, p. E04004. [7](#)
- [19] Kim, W. S., Ansar, A. I., and Steele, R. D., "Rover mast calibration, exact camera pointing, and camera handoff for visual target tracking," *ICAR '05. Proceedings., 12th International Conference on Advanced Robotics, 2005.*, 2005, pp. 384–391. doi:10.1109/ICAR.2005.1507439. [7](#)
- [20] Lowe, D. G., "Distinctive Image Features from Scale-Invariant Keypoints," *International Journal of Computer Vision*, Vol. 60, No. 2, 2004, pp. 91–110. doi:10.1023/B:VISI.0000029664.99615.94, URL <https://doi.org/10.1023/B:VISI.0000029664.99615.94>. [7](#)
- [21] Vedaldi, A., and Fulkerson, B., *Vlfeat: an open and portable library of computer vision algorithms*, an open and portable library of computer vision algorithms, ACM, New York, New York, USA, 2010. [7](#)
- [22] Yousif, K., Bab-Hadiashar, A., and Hoseinnezhad, R., "An Overview to Visual Odometry and Visual SLAM: Applications to Mobile Robotics," *Intelligent Industrial Systems*, Vol. 1, No. 4, 2015, pp. 289–311. [7](#)
- [23] Ansar, A., and Daniilidis, K., "Linear pose estimation from points or lines," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, Vol. 25, No. 5, 2003, pp. 578–589. [7](#)
- [24] Clouse, D., Padgett, C., Ansar, A., and Cheng, Y., "Automated Camera Array Fine Calibration," , 2008. [8](#)
- [25] Johnson, A., Ansar, A., Matthies, L., Trawny, N., Mourikis, A., and Roumeliotis, S., "A General Approach to Terrain Relative Navigation for Planetary Landing," *AIAA Infotech@Aerospace 2007 Conference and Exhibit*, American Institute of Aeronautics and Astronautics, Reston, Virginia, 2012, p. 391. [8](#)
- [26] Rumpler, M., Irschara, A., and Bischof, H., "Multi-View Stereo: Redundancy Benefits for 3D Reconstruction," 2011. [8](#)
- [27] Brunelli, R., *Template Matching Techniques in Computer Vision*, Theory and Practice, John Wiley & Sons, Chichester, UK, 2009. [9](#)
- [28] Tomasi, C., and Kanade, T., "Detection and Tracking of Point Features," Tech. rep., International Journal of Computer Vision, 1991. [12](#)
- [29] Myronenko, A., and Song, X. B., "Point-Set Registration: Coherent Point Drift," *CoRR*, Vol. abs/0905.2635, 2009. URL <http://arxiv.org/abs/0905.2635>. [12](#)